# Bergm: Bayesian Exponential Random Graphs in R

Alberto Caimo & Nial Friel
School of Mathematical Sciences,
University College Dublin, Ireland
{alberto.caimo,nial.friel}@ucd.ie

January 16, 2012

**Abstract**

In this paper we describe the basic features of the Bergm package for the open-source R software which provides a comprehensive framework for Bayesian analysis for exponential random graph models: tools for parameter estimation, model selection and goodness-of-fit diagnostics. We illustrate the capabilities of this package describing the algorithms that drive the package through a tutorial analysis of two well-known network datasets.

## 1 Introduction

The R package Bergm implements Bayesian analysis for Exponential Random Graph Models (ERGMs) (Wasserman and Pattison (1996); Robins et al. (2007)) using the methods described by Caimo and Friel (2011) and Caimo and Friel (2012). The package provides a comprehensive framework for Bayesian inference and model selection using Markov chain Monte Carlo (MCMC) algorithms. It can also supply graphical Bayesian goodness-of-fit procedures that address the issue of model adequacy. Although computationally intensive, the package is simple to use and represents an attractive way of analyzing network data as it offers the advantange of a complete probabilistic treatment of uncertainty. Bergm is based on the ergm package developed by Hunter et al. (2008b) and therefore it makes use of the same model set-up and network simulation

algorithms. The Bergm package has been continually improved in terms of speed performance over the last two years and one of the purposes of this paper is to highlight these improvements. We feel that this package now offers the end-user a feasible option for carrying out Bayesian inference for exponential random graphs.

Two well-known network datasets will be used throughout this tutorial for illustrative purposes: the first is the Kapferer Tailor Shop dataset (Kapferer, 1972) whose directed edges represent work interactions in a tailor shop in Zambia (then Northern Rhodesia) and nodal attributes refer to the job status. The second network is Zachary's karate club (Zachary, 1977) which represents the undirected social network graph of friendships between 34 members of a karate club at a US university in the 1970s. Figure 1 displays the graphs of these two networks.

In this paper we describe how to install and load Bergm (Section 2) providing a brief summary of what Bayesian ERGMs are (Section 3). Sections 4, 5, and 6 overview the algorithms and the functions used to produce posterior estimates for the parameters, Bayesian goodness-of-fit procedures and model selection respectively. This paper does not provide an exhaustive description of all the functionality and options available, and more information about the commands and methods mentioned are available through the R help system within the package.

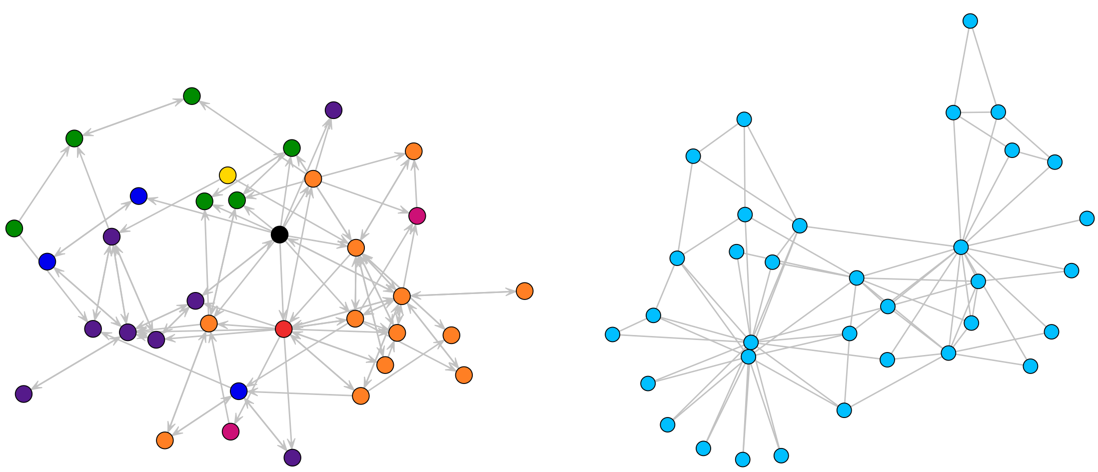


Figure 1: Kapferer Tailor Shop directed graph (right) and Zachary's karate club undirected graph (right).

## 2 Getting Bergm

The Bergm package can be obtained and loaded in R using the following commands:

```
> install.packages("Bergm")
> library("Bergm")
```

Since Bergm depends on ergm (which in turn depends on network), coda, and mvtnorm, installing the package will automatically load all the dependencies. All of these packages are available on the Comprehensive R Archive Network (CRAN) at http://CRAN.R-project.org.

The results presented in this paper have been obtained using R version 2.13.2 on a Mac using Bergm version 2.1; ergm version 2.4-3; network version 1.6; coda version 0.14-6; and mvtnorm version 0.9-999.

## 3 Bayesian exponential random graphs

The Bayesian approach to statistical problems is probabilistic. Inference is based on the posterior distribution which is the conditional probability of the unknown quantities given the observed ones. The posterior extracts the information in the data and provide a complete summary of the uncertainty about the unknowns.

In the ERGM context (see Wasserman and Pattison (1996) and Robins et al. (2007)), the purpose of Bayesian inference is to learn about the posterior distribution of the model parameters $\boldsymbol{\theta}$ of an observed graph $\boldsymbol{y}$ on $n$ nodes:

$$\pi(\boldsymbol{\theta}|\boldsymbol{y}) = \frac{p(\boldsymbol{y}|\boldsymbol{\theta})\ p(\boldsymbol{\theta})}{p(\boldsymbol{y})} = \frac{\exp\{\boldsymbol{\theta}^t s(\boldsymbol{y})\}}{z(\boldsymbol{\theta})}\ \frac{p(\boldsymbol{\theta})}{p(\boldsymbol{y})}, \tag{1}$$

where $s(\boldsymbol{y})$ is a known vector of sufficient network statistics (Figure 2) (Morris et al., 2008), $p(\boldsymbol{\theta})$ is a prior distribution placed on $\boldsymbol{\theta}$, $z(\boldsymbol{\theta})$ is the likelihood normalizing constant, and $p(\boldsymbol{y})$ is the model evidence. Equation (1) provides a probabilistic statement about how likely parameter values are after observing the data $\boldsymbol{y}$. The likelihood $p(\boldsymbol{y}|\boldsymbol{\theta})$ is translated into a proper probability distribution that can be summarised by computing expected values, standard deviations, quantiles, etc.

Unfortunately the posterior distribution (1) is doubly-intractable as both $z(\boldsymbol{\theta})$ and $p(\boldsymbol{y})$ cannot be evaluated analytically (Koskinen, 2004). This makes the use of standard MCMC procedures infeasible.

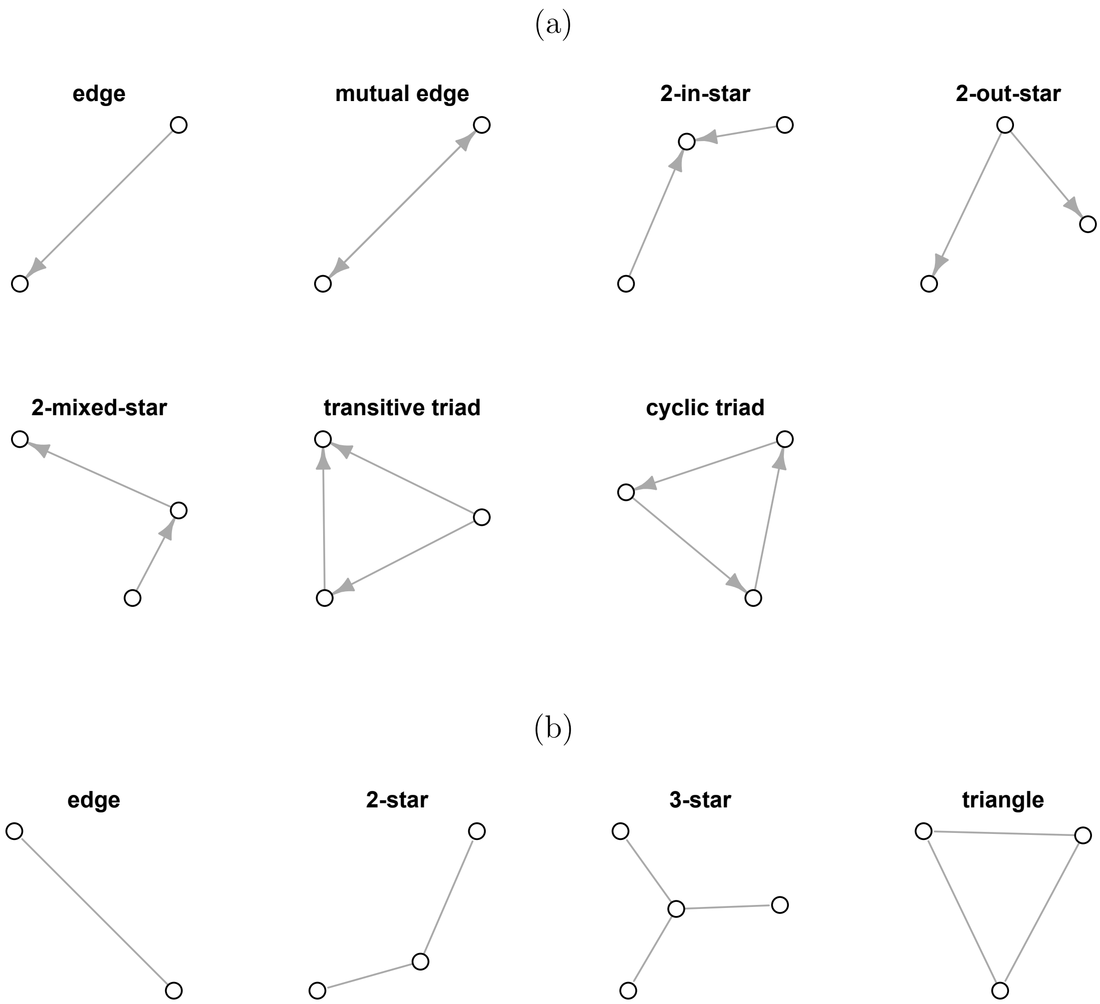



Figure 2: Some of the most used configurations for directed (a) and undirected (b) graphs.

In order to carry out Bayesian inference for ERGMs, the Bergm package makes use of a combination of Bayesian algorithms and MCMC techniques. The exchange algorithm circumvents the problem of computing the normalizing constants of the ERGM likelihoods, while the use of multiple chains interacting with each others (population MCMC approach) by means of adaptive direction sampling is able to speed up the computations and improve chain mixing quite significantly.

## 4 Bayesian parameter estimation

In order to estimate (1), the Bergm package uses the exchange algorithm described in Section 4.1 of Caimo and Friel (2011) to sample from the following distribution:

$$p(\boldsymbol{\theta}', \boldsymbol{y}', \boldsymbol{\theta}|\boldsymbol{y}) \propto p(\boldsymbol{y}|\boldsymbol{\theta})p(\boldsymbol{\theta})\epsilon(\boldsymbol{\theta}'|\boldsymbol{\theta})p(\boldsymbol{y}'|\boldsymbol{\theta}') \tag{2}$$

where $p(\boldsymbol{y}'|\boldsymbol{\theta}')$ is the likelihood on which the simulated data $\boldsymbol{y}'$ are defined, $\epsilon(\boldsymbol{\theta}'|\boldsymbol{\theta})$ is any arbitrary proposal distribution for the augmented variable $\boldsymbol{\theta}'$. As we will see in the next section, this proposal distribution is set to be a normal centered at $\boldsymbol{\theta}$.

At each MCMC iteration, the exchange algorithm consists of a Gibbs update of $\boldsymbol{\theta}'$ followed by a Gibbs update of $\boldsymbol{y}'$, which is drawn from the $p(\cdot|\boldsymbol{\theta}')$ via an MCMC algorithm (Hunter et al., 2008b). Then a deterministic exchange or swap from the current state $\boldsymbol{\theta}$ to the proposed new parameter $\boldsymbol{\theta}'$. This deterministic proposal is accepted with probability:

$$\alpha = \min\left(1, \frac{q_{\boldsymbol{\theta}}(\boldsymbol{y}')p(\boldsymbol{\theta}')h(\boldsymbol{\theta}|\boldsymbol{\theta}')q_{\boldsymbol{\theta}'}(\boldsymbol{y})}{q_{\boldsymbol{\theta}}(\boldsymbol{y})p(\boldsymbol{\theta})h(\boldsymbol{\theta}'|\boldsymbol{\theta})q_{\boldsymbol{\theta}'}(\boldsymbol{y}')} \times \frac{z(\boldsymbol{\theta})z(\boldsymbol{\theta}')}{z(\boldsymbol{\theta})z(\boldsymbol{\theta}')}\right),$$

where $q_{\boldsymbol{\theta}}$ and $q_{\boldsymbol{\theta}'}$ indicates the unnormalised likelihoods with parameter $\boldsymbol{\theta}$ and $\boldsymbol{\theta}'$, respectively. Notice that all the normalising constants cancel above and below in the fraction above, in this way avoiding the need to calculate the intractable normalising constant.

The exchange algorithm is implemented by the `bergm` function in the following way:

---

`for` $i = 1, \ldots, N$

1. `generate` $\boldsymbol{\theta}'$ `from` $\epsilon(\cdot|\boldsymbol{\theta})$

2. `simulate` $\boldsymbol{y'}$ `from` $p(\cdot|\boldsymbol{\theta}')$

3. `update` $\boldsymbol{\theta} \leftarrow \boldsymbol{\theta}'$ `with probability`

$$\log(\alpha) = \min\left(0,\ [\boldsymbol{\theta} - \boldsymbol{\theta}']^t\,[s(\boldsymbol{y'}) - s(\boldsymbol{y})] + \log\left[\frac{p(\boldsymbol{\theta}')}{p(\boldsymbol{\theta})}\right]\right) \tag{3}$$

`end for`

---

where $s(\boldsymbol{y})$ is the observed vector of network statistics and $s(\boldsymbol{y'})$ is the simulated vector of network statistics.

## 4.1 Block-update sampler

Step 1 of the algorithm consists in generating $\boldsymbol{\theta}'$ from some proposal distribution within each iteration. Bergm uses a block-update sampler with normal proposal to simultaneously update of the parameter values in the MCMC chain:

$$\epsilon(\boldsymbol{\theta}'|\boldsymbol{\theta}) \sim \mathcal{N}(\boldsymbol{\theta}, \boldsymbol{\Sigma}_{\boldsymbol{\epsilon}}). \tag{4}$$

Typically, tuning the proposal distribution $\boldsymbol{\epsilon}$ from which $\boldsymbol{\theta}'$ is drawn represents the crucial part of the algorithm since a poor tuning of the proposal parameter $\boldsymbol{\Sigma}_{\boldsymbol{\epsilon}}$ can slow down the chain's mixing rate and therefore the algorithm can take a very long time to converge to the stationary posterior density.

## 4.2 Parallel adaptive direction sampler

In order to improve mixing a parallel adaptive direction sampler (ADS) is considered: at each $i$-th iteration of the algorithm we have a collection of $H$ different chains interacting with one another. By construction, the state space consists of $\{\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_H\}$ with target distribution $p(\boldsymbol{\theta}_1|\boldsymbol{y}) \otimes \cdots \otimes p(\boldsymbol{\theta}_H|\boldsymbol{y})$. A parallel ADS move consists of generating a new value $\boldsymbol{\theta}'_h$ from the difference of two parameters $\boldsymbol{\theta}_{h_1}$ and $\boldsymbol{\theta}_{h_2}$ (randomly selected from other chains) multiplied by a scalar term $\gamma$ which is called parallel ADS move factor plus a random term $\boldsymbol{\epsilon}$ called parallel ADS move parameter (Figure 3) which is equivalent to the block-update sampler defined in (4). The algorithm can be summarised as follows:

```
for i = 1,...,N
    for h = 1,...,H
        1.  generate h_1 and h_2 such that h_1 ≠ h_2 ≠ h
        2.  generate θ'_h from γ(θ_{h_1} − θ_{h_2}) + ε(·|θ_h)
        3.  simulate y' from p(·|θ'_h)
        4.  update θ_h ← θ'_h with probability
```

$$\min\left(0, [\boldsymbol{\theta}_h - \boldsymbol{\theta}'_h]^t \left[s(\boldsymbol{y}') - s(\boldsymbol{y})\right] + \log\left[\frac{p(\boldsymbol{\theta}'_h)}{p(\boldsymbol{\theta}_h)}\right]\right) \tag{5}$$

```
    end for
end for
```

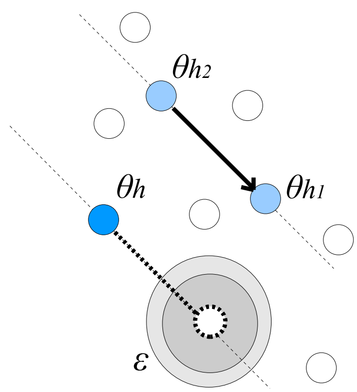



Figure 3: The parallel ADS move of the current state (darker blue dot) consists of generating a new parameter value along the direction made by the difference of two randomly sampled parameter states (light blue dots) belonging to different chains plus a random term $\boldsymbol{\epsilon}$.

### 4.3 Kapferer tailor shop network

Consider the Kapferer Tailor Shop network and a 3-dimensional model including the following network statistics: edges (`edges`), mutual edges (`mutual`) and cyclic triples (`ctriple`) involving nodes with the same job status, where the job status is represented by a categorical variable of 8 levels:

| | |
|---|---|
| `edges` | $\sum_{i<j} y_{ij}$ |
| `mutual` | $\sum_{i<j} y_{ij}y_{ji}$ |
| `ctriple("job")` | $\sum_{i<j<k} y_{ij}y_{jk}y_{ki}$ where $i, j, k$ have the same job status |

The format of the model specification is the same of an `ergm` formula:

```
> formula <- y ~ edges + mutual + ctriple("job")
```

Then we can use the `bergm` function to sample from the posterior distribution using the MCMC algorithm described above. In this example we use the parallel ADS procedure. By default, the number of chains is set as twice the number of dimensions of the model. It is possible to choose a different number of chains by using the argument `nchains`. In order to perform the block-site update described in Section 4.1 it is necessary to set `nchains` = 1. For each chain, we can then set the number of burn-in iterations (`burn.in`) and the number of iterations after the burn-in (`main.iters`). The number of iterations used to simulate a network $\boldsymbol{y}'$ at each iteration is defined by the argument `aux.iters`.

```
> post.est  <- bergm(formula,
+                    burn.in=200,
+                    gamma=0.7,
+                    main.iters=1000,
+                    aux.iters=25000)
```

The population MCMC with parallel ADS move is the default procedure of the `bergm` function. The total number of iterations (eg the size of the posterior sample) is `nchains` $\times$ `main.iters`. The proposal covariance structure $\boldsymbol{\Sigma}_\epsilon$ is defined by the argument `sigma.epsilon` which is set to be a diagonal matrix with every diagonal entry equal to a small number. In many cases, good mixing of the chain is ensured by a sensible tuning of the parallel ADS move factor `gamma` and therefore the argument `sigma.epsilon` can

be generally left at its default setting. As said above, parallel ADS is adopted as the default procedure but it is automatically disabled in the case of uni-dimensional models where the block-update sampler is used and the argument `gamma` is used to tune the variance of the normal proposal distribution $\epsilon$.

After completing the estimation, `post.est` is an object of the class `bergm` and contains a list of attributes among which are the real and CPU time (in seconds) taken by the estimation process:

```
> post.est$time
   user  system elapsed
 99.083   0.759 100.659
```

It is possible to visualise the results of the MCMC estimation by using the `bergm.output` function which is based on the `coda` package (Plummer et al., 2006) which is an R package for MCMC output analysis and diagnostics.

```
> bergm.output(post.est,lag.max=50)

 MCMC results for Model: y ~ edges + mutual + ctriple("job")


 Posterior mean:
          theta1 (edges) theta2 (mutual) theta3 (ctriple.job)
Chain 1       -3.4256853       3.7647452            0.8168465
Chain 2       -3.4091174       3.7447376            0.8567871
Chain 3       -3.4140165       3.7377072            0.8532598
Chain 4       -3.4366544       3.7868822            0.8585445
Chain 5       -3.4201932       3.8034972            0.8340857
Chain 6       -3.4009597       3.7021253            0.8433665


 Posterior sd:
          theta1 (edges) theta2 (mutual) theta3 (ctriple.job)
Chain 1        0.2046574       0.4287769            0.1882377
Chain 2        0.2291025       0.4960458            0.1829550
Chain 3        0.1903504       0.4453100            0.1700874
Chain 4        0.2673414       0.5266250            0.1655076
```

```
Chain 5       0.2474045      0.4652013            0.1543426
Chain 6       0.2696806      0.4775102            0.1949814

          Acceptance rate:
Chain 1             0.213
Chain 2             0.267
Chain 3             0.246
Chain 4             0.224
Chain 5             0.260
Chain 6             0.237


 Overall posterior density estimate:
           theta1 (edges) theta2 (mutual) theta3 (ctriple.job)
Post. mean     -3.4177711       3.7499491            0.8438150
Post. sd        0.2373836       0.4768709            0.1782518

 Overall acceptance rate: 0.241166666666667
```

The output above shows the results of the MCMC estimation: posterior means and standard deviations, and acceptance rates for every chain in the population and for the overall chain. Notice that the posterior summaries of each chain are consistent with each other. Figure 4 displays the MCMC diagnostic plots produced by the `bergm.output` function. In this example, we notice a low density effect expressed by the negative value of the posterior mean of the density effect parameter (edges) combined with the positive mutuality and transitivity within people having the same job status expressed by the mutual edge and cyclic triple parameters respectively. The overall acceptance rate is around 24% and the autocorrelation is negligible after lag 50.

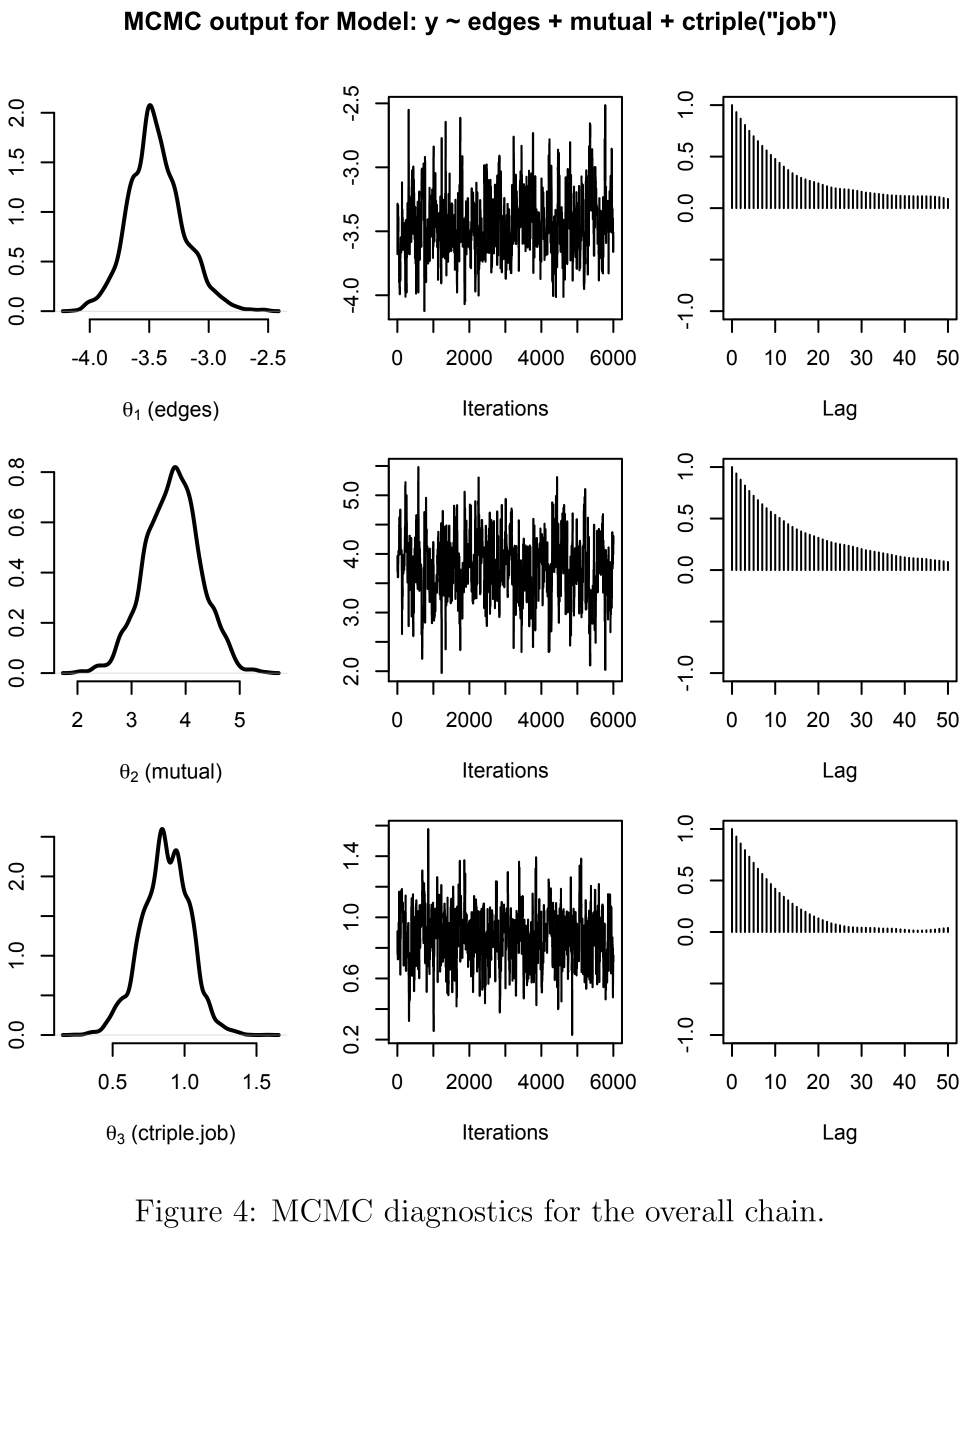



Figure 4: MCMC diagnostics for the overall chain.

# 5 Bayesian goodness-of-fit diagnostics

The `bgof` function provides a useful tool for assessing Bayesian goodness-of-fit so as to examine the fit of the data to the posterior model obtained by the `bergm` function. The observed network data $\boldsymbol{y}$ is compared with a set of networks $\boldsymbol{y}_1, \boldsymbol{y}_2, \ldots, \boldsymbol{y}_S$ simulated from $S$ independent realisations $\boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \ldots, \boldsymbol{\theta}_S$ of the posterior density estimate. This comparison is made in terms of high-level characteristics $g(\cdot)$ such as higher degree distributions, etc. (see Hunter et al. (2008a)). The algorithm can be summarised as follows:

---

for $i = 1, \ldots, S$

1. sample $\boldsymbol{\theta}_i$ from the estimate of $p(\boldsymbol{\theta}|\boldsymbol{y})$

3. simulate $\boldsymbol{y}_i$ from $p(\cdot|\boldsymbol{\theta}_i)$

4. calculate $g(\boldsymbol{y}_i)$

end for

---

For example, the code below is used to compare the Kapferer Tailor Shop network with a series of networks simulated from 100 random realisations (`sample.size`) of the estimated posterior distribution `post.est` using $50,000$ iterations (`aux.iters`) for the network simulation step. The `bgof` function may take a few seconds to run and, at the end of the execution, it will automatically plot the results as shown in Figure 5.

```
> bgof(post.est,
+       sample.size=100,
+       aux.iters=50000,
+       directed=TRUE,
+       n.ideg=20,
+       n.odeg=20,
+       n.dist=10,
+       n.esp=15)
```

The set of statistics used for the comparison of directed networks includes the in-degree

distribution, the out-degree distribution, the minimum geodesic distance distribution and the edgewise shared partner distribution. The arguments `n.ideg`, `n.odeg`, `n.dist`, and `n.esp` indicates the number of boxplots to plot for each distribution respectively.

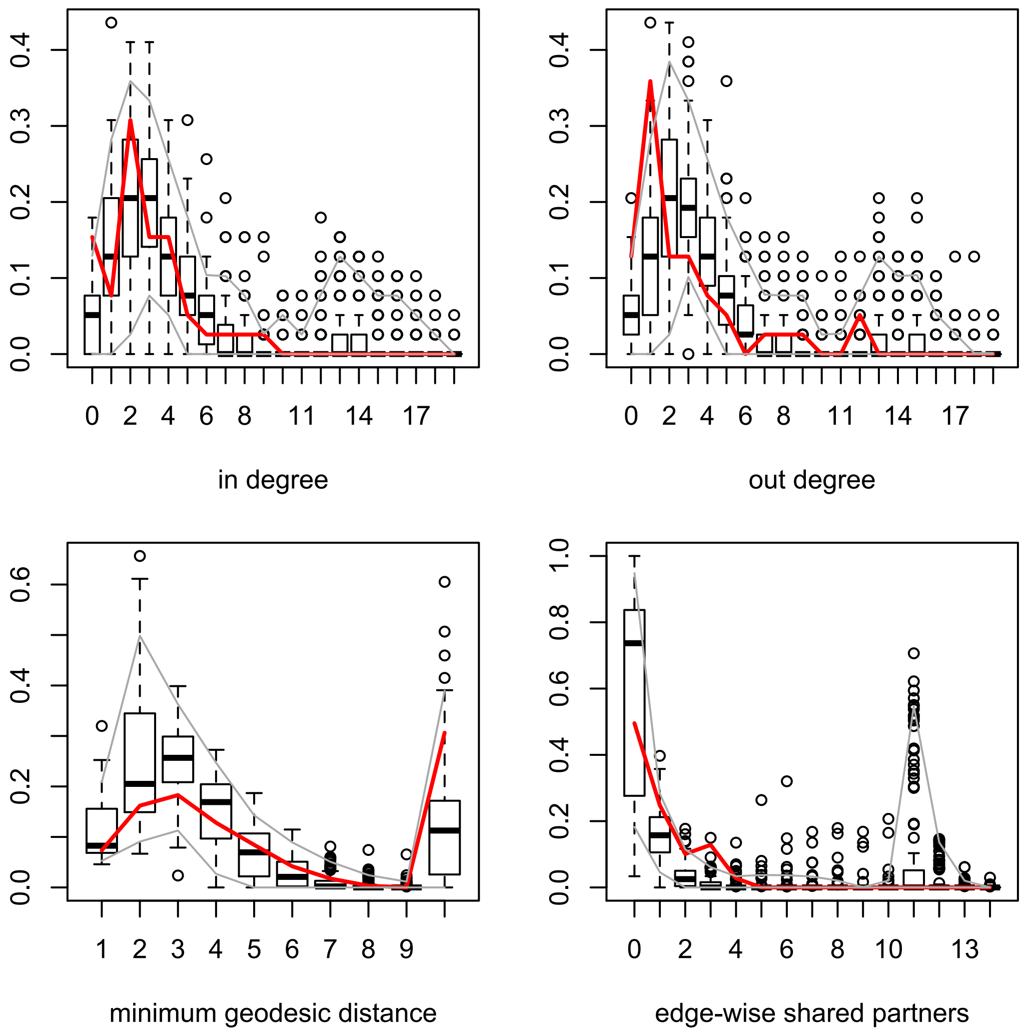



Figure 5: Bayesian goodness-of-fit diagnostics. The red line displays the goodness of fit statistics for the observed data together with boxplots of goodness of fit statistics based on 100 simulated networks from the posterior distribution.

In Figure 5 we see, based on the various goodness of fit statistics, that the networks simulated from the posterior distribution are in reasonable agreement with the observed network. We can therefore conclude that the data is a reasonable fit to the model, despite its simplicity.

# 6 Bayesian model selection

An important problem in statistical analysis is the choice of an optimal model from a set of *a priori* competing models. In the ERGM context, this task translates into the choice of which subset of network statistics should be included into the model.

Let $m$ indicate a particular model from a set of competing models with corresponding parameters $\boldsymbol{\theta}$. Following the Bayesian paradigm, interest focuses on exploring the posterior distribution,

$$p(m, \boldsymbol{\theta}|\boldsymbol{y}) \propto p(\boldsymbol{y}|m, \boldsymbol{\theta})\; p(\boldsymbol{\theta}|m)\; p(m) \tag{6}$$

where $p(\boldsymbol{\theta}|m)$ and $p(m)$ are prior distributions within model $m$, and on model $m$, respectively. The reversible Jump Markov chain Monte Carlo (RJMCMC) algorithm (Green, 1995) was designed to explore this type of posterior distribution across the joint model and parameter space. It is therefore a type of MCMC algorithm that allows one to jointly explore the uncertain between and within models.

This approach is very appealing since it relies exclusively on probabilistic considerations but is very challenging from a computational viewpoint. As stated above, the intractability of the likelihood normalising constant $z(\boldsymbol{\theta})$ in (1) renders standard RJMCMC techniques infeasible. However, the exchange algorithm used for parameter estimation can be easily generalised so as to include model indicators.

The auto-RJ exchange algorithm described in Caimo and Friel (2012) represents a trans-dimensional RJMCMC extension of the exchange algorithm involving an independence sampler based on a distribution fitting a parametric density approximation to the within-model posterior. This approach overcomes the issue of the likelihood intractability sampling from:

$$p(\boldsymbol{\theta}', \boldsymbol{\theta}, m', m, \boldsymbol{y}'|\boldsymbol{y}) \propto p(\boldsymbol{y}|\boldsymbol{\theta}, m)p(\boldsymbol{\theta}|m)p(m)w(\boldsymbol{\theta}'|m')h(m'|m)p(\boldsymbol{y}'|\boldsymbol{\theta}', m') \tag{7}$$

where $m$ and $m'$ are two competing models, $p(\boldsymbol{y}|\boldsymbol{\theta}, m)$ and $p(\boldsymbol{y}'|\boldsymbol{\theta}', m')$ are the two likelihoods for the data $\boldsymbol{y}$ under model $m$ and the simulated data $\boldsymbol{y}'$ under model $m'$ respectively, $p(\boldsymbol{\theta}|m)$ and $p(m)$ are the priors for the parameter and the respective model $m$, $w(\boldsymbol{\theta}'|m')$ is a within-model proposal (independence sampler) which fit a parametric density approximation to the model posteriors and $h(m'|m)$ is a between-model proposal. Notice that the marginal distribution for for $\boldsymbol{\theta}'$ and $m'$ in (7) is the target distribution of interest $p(\boldsymbol{\theta}, m|\boldsymbol{y})$.

The `bergmS` function implements the auto-RJ exchange algorithm which consists of two parts: an offline step and an online step. In the first step, samples from the posterior $p(\boldsymbol{\theta}|\boldsymbol{y}, m)$ are gathered from each competing model using the `bergm` function and then approximated by normal distributions $\mathcal{N}(\hat{\boldsymbol{\mu}}_m, \hat{\boldsymbol{\Sigma}}_m)$ determined by the first and second moments from a sample from the model.

The second step (online run) of the algorithm consists of a Gibbs update of $m'$ followed by a Gibbs update of $\boldsymbol{\theta}'$ which is generated via the independence sampler $w(\boldsymbol{\theta}'|m') \sim \mathcal{N}(\hat{\boldsymbol{\mu}}_{m'}, \hat{\boldsymbol{\Sigma}}_{m'})$. This is followed by a Gibbs update of $\boldsymbol{y}'$ which is generated from $p(\cdot|\boldsymbol{\theta}', m')$. Then a deterministic exchange move from a current state $(\boldsymbol{\theta}, m)$ to the proposed new state $(\boldsymbol{\theta}', m')$ is accepted with probability:

$$\alpha = \min\left\{1, \frac{q_{\boldsymbol{\theta},m}(\boldsymbol{y}')}{q_{\boldsymbol{\theta},m}(\boldsymbol{y})} \frac{q_{\boldsymbol{\theta}',m'}(\boldsymbol{y})}{q_{\boldsymbol{\theta}',m'}(\boldsymbol{y}')} \frac{p(\boldsymbol{\theta}'|m')}{p(\boldsymbol{\theta}|m)} \frac{p(m')}{p(m)} \frac{w(\boldsymbol{\theta}|\hat{\boldsymbol{\mu}}_m, \hat{\boldsymbol{\Sigma}}_m)}{w(\boldsymbol{\theta}'|\hat{\boldsymbol{\mu}}_{m'}, \hat{\boldsymbol{\Sigma}}_{m'})}\right\}. \tag{8}$$

where $q_{\boldsymbol{\theta},m}$ and $q_{\boldsymbol{\theta}',m'}$ indicates the unnormalised likelihoods under model $m$ with parameter $\boldsymbol{\theta}$ and under model $m'$ with parameter $\boldsymbol{\theta}'$ respectively.

The structure of the `bergmS` function can be described in the following way:

---

`for` $i = 1, \ldots, N$

1. `generate` $m'$ from $h(\cdot|m)$
2. `generate` $\boldsymbol{\theta}' \sim \mathcal{N}(\hat{\boldsymbol{\mu}}_{m'}, \hat{\boldsymbol{\Sigma}}_{m'})$
3. `simulate` $\boldsymbol{y}'$ `from` $p(\cdot|\boldsymbol{\theta}', m')$
4. `update` $(\boldsymbol{\theta}, m) \leftarrow (\boldsymbol{\theta}', m')$ `with probability:`

$$\log(\alpha) = \min\left(0, \boldsymbol{\theta}^t[s_m(\boldsymbol{y}') - s_m(\boldsymbol{y})] + \boldsymbol{\theta}'^t[s_{m'}(\boldsymbol{y}') - s_{m'}(\boldsymbol{y})] + \log\left[\frac{p(\boldsymbol{\theta}')}{p(\boldsymbol{\theta})} \frac{w(\boldsymbol{\theta}|\hat{\boldsymbol{\mu}}_m, \hat{\boldsymbol{\Sigma}}_m)}{w(\boldsymbol{\theta}'|\hat{\boldsymbol{\mu}}_{m'}, \hat{\boldsymbol{\Sigma}}_{m'})}\right]\right)$$

`end for`

---

where $s_m(\boldsymbol{y})$ and $s_{m'}(\boldsymbol{y})$ are the observed vectors of network statistics under model $m$ and $m'$ respectively, and $s_m(\boldsymbol{y}')$ and $s_{m'}(\boldsymbol{y}')$ are the simulated vector of network statistics under model $m$ and $m'$ respectively. The reader is referred to Caimo and Friel (2012) for much details on this algorithm.

### 6.1 Karate club network

Consider the Karate club network, we propose three models to fit the data using a set of new specification statistics introduced by Snijders et al. (2006): geometrically weighted edgewise shared partners (`gwesp`) and geometrically weighted degrees (`gwdegree`):

$$\begin{array}{ll} \texttt{gwesp} & e^{\phi_v} \sum_{k=1}^{n-2} \left\{ 1 - \left(1 - e^{-\phi_v}\right)^k \right\} EP_k(\boldsymbol{y}) \\ \texttt{gwdegree} & e^{\phi_u} \sum_{k=1}^{n-1} \left\{ 1 - \left(1 - e^{-\phi_u}\right)^k \right\} D_k(\boldsymbol{y}) \end{array}$$

where the scale parameters $\phi_v = 0.2$ and $\phi_u = 0.8$. $D_k(\boldsymbol{y})$ is the number of pairs that have exactly $k$ common neighbours and $EP_k(\boldsymbol{y})$ is the number of connected pairs with exactly $k$ common neighbours. The specification of these models requires the creation of a list of formulas:

```
> formulae <- c(y ~ edges + gwesp(0.2,fixed=TRUE),
+                y ~ edges + gwdegree(0.8,fixed=TRUE),
+                y ~ edges + gwesp(0.2,fixed=TRUE)
+                          + gwdegree(0.8,fixed=TRUE))
```

The `bergmS` command is then used to carry out the algorithm. To do this we have to specify several arguments for both the offline and online step.

The offline run consists of running the `bergm` function for each of the models proposed. Therefore we set some arguments `main.iters`, `burn.ins`, `gammas` which are vectors containing values for `bergm` arguments: `main.iters`, `burn.in`, `gamma` for each competing model.

The argument `iters` refers to the number of iterations used for the online run. The number of MCMC steps used for network simulation is specified as usual by the argument `aux.iters` and this will be used in both the offline and the online step. The command below should take around 10 minutes depending on the CPU speed of the computer.

```
> mod.sel <- bergmS(formulae,
+                   iters=25000,
+                   aux.iters=15000,
+                   main.iters=rep(700,3),
+                   burn.ins=rep(100,3),
+                   gammas=c(1,1,0.8))
```

The `bergmS.output` function produces the MCMC diagnostics for each competing model explored by the MCMC algorithm. Figure 6 and 7 displays the plots regarding the posterior model and parameter density estimate respectively.

```
> best.mod <- bergmS.output(mod.sel)

 BEST MODEL
 ----------
Model 1: y ~ edges + gwesp(0.2, fixed = TRUE)

 Posterior parameter estimate:
                          Post. mean:  Post. sd:
theta1 (edges)             -3.2574625  0.3278196
theta2 (gwesp.fixed.0.2)    1.1008261  0.2515162

 Within-model acceptance rate: 0.26


Model 3: y ~ edges + gwesp(0.2, fixed = TRUE) + gwdegree(0.8, fixed = TRUE)

 Posterior parameter estimate:
                          Post. mean:  Post. sd:
theta1 (edges)             -3.4916584  0.5146103
theta2 (gwesp.fixed.0.2)    1.1824831  0.2737858
theta3 (gwdegree)           0.4783638  0.5852124

 Within-model acceptance rate: 0.14

BF_13 = 13.4508670520231

Between-model acceptance rate: 0.03
```

In the results above we have the posterior parameter estimates for two of the competing models (Model 2 has not been visited through the MCMC runs) with respective

within-model acceptance rates and an estimate of the Bayes Factor (about 13) for the comparison between Model 1 and Model 3 which makes clear that there is evidence that Model 1 is the best model of the set.

After running the command `bergmS.output`, it is possible to perform a Bayesian goodness-of-fit tests. In this case, since the observed network is undirected, the set of high-level statistics include the degree distribution in place of the in-degree and out-degree distributions.

```
> bgof(best.mod,
+       aux.iters=30000,
+       n.deg=20,
+       n.dist=10,
+       n.esp=15)
```

In this example, the observed data appears to be a reasonable fit to the posterior distribution of the model selected (Model 1), based on the goodness of fit geodesic and the shared partners statistics displayed in Figure 8.

## 7 Discussion

The software package Bergm aims to help researchers and practitioners in two ways. Firstly, it is currently the only package for R that provides a simple and complete range of tools for conducting Bayesian analysis for exponential random graph models. Secondly, Bergm makes available a platform that can be easily customised, extended, and adapted to address different requirements.

The software package is under continual development and it is far from finished. The main limitation of the software is its computational cost which makes it unsuitable for managing network graphs larger than hundreds of nodes, however it is perfectly suited for networks involving up to a hundred nodes. However an important improvement in terms of computational time and efficiency will be done by turning some of the R functions into C functions integrated with the ergm package. We expect that will yield further reductions in computional run time.

Future versions of the Bergm package will address several issues including Bayesian analysis of Curved Exponential Random Graph Models (Hunter and Handcock, 2006)

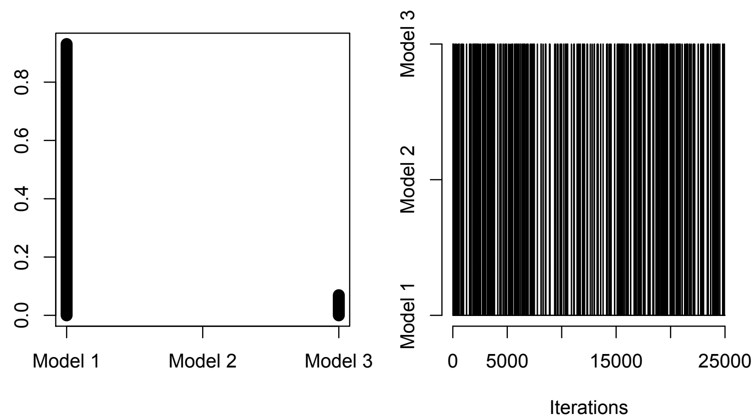



Figure 6: MCMC diagnostics: posterior model probabilities.

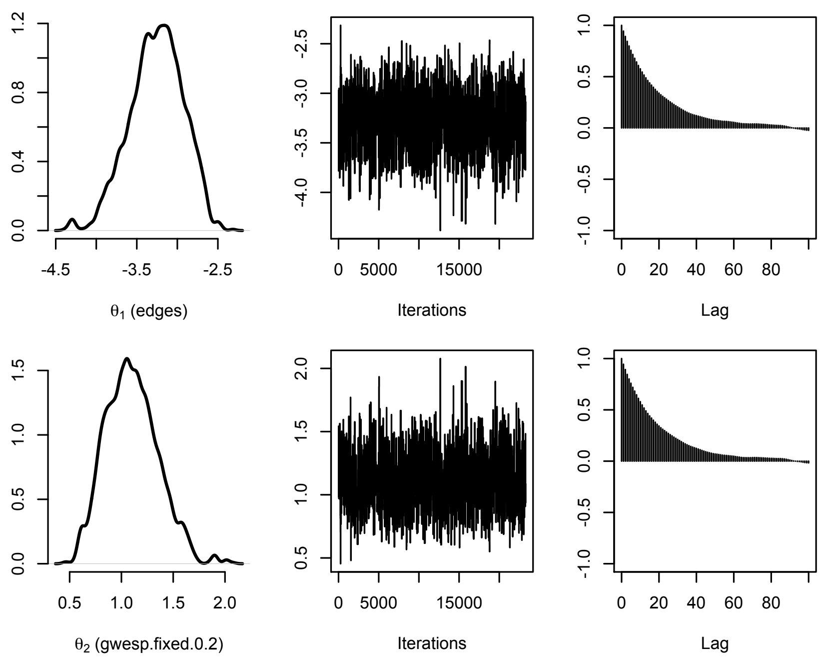



Figure 7: MCMC diagnostics: posterior parameter probabilities.

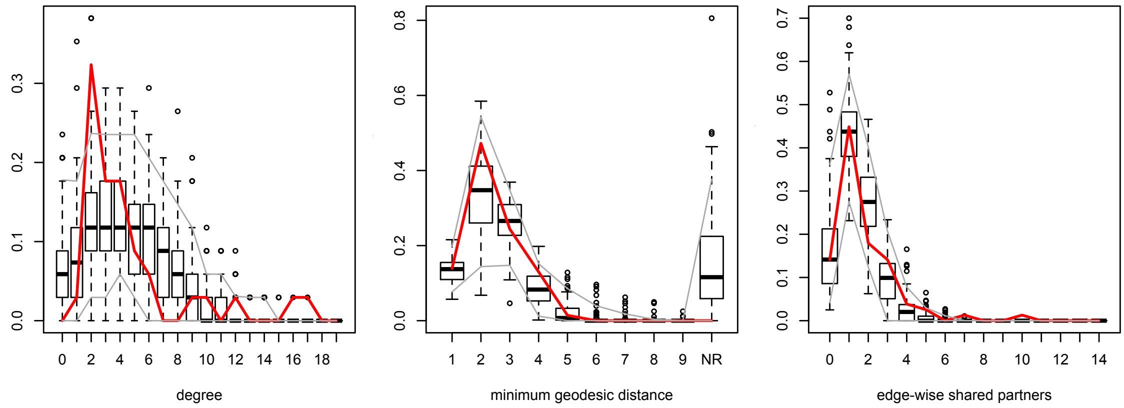



Figure 8: Bayesian goodness-of-fit diagnostics.

and exponential random graph models with missing data.

**Acknowledgement** Alberto Caimo was supported by an IRCSET Embark Initiative award and Nial Friel's research was supported by a Science Foundation Ireland Research Frontiers Program grant, 09/RFP/MTH2199.

# References

Caimo, A. and Friel, N. (2011), "Bayesian inference for exponential random graph models," *Social Networks*, 33, 41 – 55.

— (2012), "Bayesian model selection for exponential random graph models," Tech. rep., University College Dublin, available in e-print format at http://arxiv.org/abs/1201.2337.

Green, P. J. (1995), "Reversible jump Markov chain Monte Carlo computation and Bayesian model determination," *Biometrika*, 82, 711–732.

Hunter, D. R., Goodreau, S. M., and Handcock, M. S. (2008a), "Goodness of Fit of Social Network Models," *Journal of the American Statistical Association*, 103, 248–258.

Hunter, D. R. and Handcock, M. S. (2006), "Inference in curved exponential family

models for networks," *Journal of Computational and Graphical Statistics*, 15, 565–583.

Hunter, D. R., Handcock, M. S., Butts, C. T., Goodreau, S. M., and Morris, M. (2008b), "ergm: A Package to Fit, Simulate and Diagnose Exponential-Family Models for Networks," *Journal of Statistical Software*, 24, 1–29.

Kapferer, B. (1972), *Strategy and transaction in an African factory: African workers and Indian management in a Zambian town*, no. 10, Manchester University Press.

Koskinen, J. H. (2004), "Bayesian Analysis of Exponential Random Graphs - Estimation of Parameters and Model Selection," *Research Report 2004:2, Department of Statistics, Stockholm University.*

Morris, M., Handcock, M. S., and Hunter, D. R. (2008), "Specification of Exponential-Family Random Graph Models: Terms and Computational Aspects," *Journal of Statistical Software*, 24.

Plummer, M., Best, N., Cowles, K., and Vines, K. (2006), "CODA: Convergence Diagnosis and Output Analysis for MCMC," *R News*, 6, 7–11.

Robins, G., Pattison, P., Kalish, Y., and Lusher, D. (2007), "An introduction to exponential random graph models for social networks," *Social Networks*, 29, 169–348.

Snijders, T. A. B., Pattison, P. E., Robins, G. L., and S., H. M. (2006), "New specifications for exponential random graph models," *Sociological Methodology*, 36, 99–153.

Wasserman, S. and Pattison, P. (1996), "Logit models and logistic regression for social networks: I. An introduction to Markov graphs and $p^*$," *Psycometrica*, 61, 401–425.

Zachary, W. (1977), "An information flow model for conflict and fission in small groups," *Journal of Anthropological Research*, 33, 452–473.